

CAPACITY AND TRAINABILITY IN RECURRENT NEURAL NETWORKS

Jasmine Collins*, Jascha Sohl-Dickstein & David Sussillo

Google Brain

Google Inc.

Mountain View, CA 94043, USA

{jllcollins, jaschasd, sussillo}@google.com

ABSTRACT

Two potential bottlenecks on the expressiveness of recurrent neural networks (RNNs) are their ability to store information about the task in their parameters, and to store information about the input history in their units. We show experimentally that all common RNN architectures achieve nearly the same per-task and per-unit capacity bounds with careful training, for a variety of tasks and stacking depths. They can store an amount of task information which is linear in the number of parameters, and is approximately 5 bits per parameter. They can additionally store approximately one real number from their input history per hidden unit. We further find that for several tasks it is the per-task parameter capacity bound that determines performance. These results suggest that many previous results comparing RNN architectures are driven primarily by differences in training effectiveness, rather than differences in capacity. Supporting this observation, we compare training difficulty for several architectures, and show that vanilla RNNs are far more difficult to train, yet have higher capacity. Finally, we propose two novel RNN architectures, one of which is easier to train than the LSTM or GRU.

1 INTRODUCTION

Research and application of recurrent neural networks (RNNs) have seen explosive growth over the last few years, (Martens & Sutskever, 2011; Graves et al., 2009), and RNNs have become the central component for some very successful model classes and application domains in deep learning (speech recognition (Amodei et al., 2015), seq2seq (Sutskever et al., 2014), neural machine translation (Bahdanau et al., 2014), the DRAW model (Gregor et al., 2015), educational applications (Piech et al., 2015), and scientific discovery (Mante et al., 2013)). Despite these recent successes, it is widely acknowledged that designing and training the RNN components in complex models can be extremely tricky. Painfully acquired RNN expertise is still crucial to the success of most projects.

One of the main strategies involved in the deployment of RNN models is the use of the Long Short Term Memory (LSTM) networks (Hochreiter & Schmidhuber, 1997), and more recently the Gated Recurrent Unit (GRU) proposed by Cho et al. (2014); Chung et al. (2014) (we refer to these as gated architectures). The resulting models are perceived as training more easily, and achieving lower error. While it is widely appreciated that RNNs are universal approximators (Doya, 1993), an unresolved question is the degree to which gated models are more computationally powerful in practice, as opposed to simply being easier to train.

Here we provide evidence that the observed superiority of gated models over vanilla RNN models is almost exclusively driven by trainability. First we describe two types of capacity bottlenecks that various RNN architectures might be expected to suffer from: parameter efficiency related to learning the task, and the ability to remember input history. Next, we describe our experimental setup where we disentangle the effects of these two bottlenecks, including training with extremely thorough hyperparameter (HP) optimization. Finally, we describe our capacity experiment results (per-parameter and per-unit), as well as the results of trainability experiments (training on extremely hard tasks where gated models might reasonably be expected to perform better).

*Work done as a member of the Google Brain Residency program (g.co/brainresidency)

1.1 CAPACITY BOTTLENECKS

There are several potential bottlenecks for RNNs, for example: How much information about the task can they store in their parameters? How much information about the input history can they store in their units? These first two bottlenecks can both be seen as memory capacities (one for the task, one for the inputs), for different types of memory.

Another, different kind of capacity stems from the set of computational primitives an RNN is able to perform. For example, maybe one wants to multiply two numbers. In terms of number of units and time steps, this task may be very straight-forward using some specific computational primitives and dynamics, but with others it may be extremely resource heavy. One might expect that differences in computational capacity due to different computational primitives would play a large role in performance. However, despite the fact that the gated architectures are outfitted with a multiplicative primitive between hidden units, while the vanilla RNN is not, we found no evidence of a computational bottleneck in our experiments. We therefore will focus only on the per-parameter capacity of an RNN to learn about its task during training, and on the per-unit memory capacity of an RNN to remember its inputs.

1.2 EXPERIMENTAL SETUP

RNNs have many HPs, such as the scalings of matrices and biases, and the functional form of certain nonlinearities. There are additionally many HPs involved in training, such as the choice of optimizer, and the learning rate schedule. In order to train our models we employed a HP tuner that uses a Gaussian Process model similar to Spearmint (see Appendix, section on HP tuning and Desautels et al. (2014); Snoek et al. (2012) for related work). The basic idea is that one requests HP values from the tuner, runs the optimization to completion using those values, and then returns the validation loss. This loss is then used by the tuner, in combination with previously reported losses, to choose new HP values such that over many experiments, the validation loss is minimized with respect to the HPs. For all of our experiments, we report the evaluation loss (separate from the validation loss returned to the HP optimizer) after the HP tuner has highly optimized the task (hundreds to many thousands of experiments for each architecture and task).

In our studies we used a variety of well-known RNN architectures: standard RNNs such as the vanilla RNN and the newer IRNN (Le et al., 2015), as well as gated RNN architectures such as the GRU and LSTM. We rounded out our set of models by innovating two novel (to our knowledge) RNN architectures (see Section 1.4) we call the Update Gate RNN (UGRNN), and the Intersection RNN (+RNN). The UGRNN is a ‘minimally gated’ RNN architecture that has only a coupled gate between the recurrent hidden state, and the update to the hidden state. The +RNN uses coupled gates to gate both the recurrent and depth dimensions in a straightforward way.

To further explore the various strengths and weaknesses of each RNN architecture, we also used a variety of network depths: 1, 2, 4, 8, in our experiments.¹ In most experiments, we held the number of parameters fixed across different architectures and different depths. More precisely, for a given experiment, a maximum number of parameters was set, along with an input and output dimension. The number of hidden units per layer was then chosen such that the number of parameters, summed across all layers of the network, was as large as possible without exceeding the allowed maximum.

For each of our 6 tasks, 6 RNN variants, 4 depths, and 6+ model sizes, we ran the HP tuner in order to optimize the relevant loss function. Typically this resulted in many hundreds to several thousands of HP evaluations, each of which was a full training run up to millions of training steps. Taken together, this amounted to CPU-millennia worth of computation.

1.3 RELATED WORK

While it is well known that RNNs are universal approximators of arbitrary dynamical systems (Doya, 1993), there is little theoretical work on the task-capacity of RNNs. Koiran & Sontag (1998) studied the VC dimension of RNNs, which provides an upper bound on their task-capacity (defined in Section 2.1). These upper bounds are not a close match to our experimental results. For instance, we find

¹Not all experiments used a depth of 8, due to limits on computational resources.

that performance saturates rapidly in terms of the number of unrolling steps (Figure 2b), while the relevant bound increases linearly with the number of unrolling steps.

Empirically, Karpathy et al. (2015) have studied how LSTMs encode information in character-based text modeling tasks. Further, Sussillo & Barak (2013) have reverse-engineered the vanilla RNN trained on simple tasks, using the tools and language of nonlinear dynamical systems theory. In Foerster et al. (2016) the behavior of switched affine recurrent networks is carefully examined.

The ability of RNNs to store information about their input has been better studied, in both the context of machine learning and theoretical neuroscience. Previous work on short term memory traces explores the tradeoffs between memory fidelity and duration, for the case that a new input is presented to the RNN at every time step (Jaeger & Haas, 2004; Maass et al., 2002; White et al., 2004; Ganguli et al., 2008; Charles et al., 2014). We use a simpler capacity measure consisting only of the ability of an RNN to store a single input vector after a single presentation. Our results suggest that, contrary to common belief, the capacity of RNNs to remember their input history is not a practical limiting factor on their performance.

The precise details of what makes an RNN architecture perform well is an extremely active research field (e.g. Jozefowicz et al. (2015)). A highly related article is Greff et al. (2015), in which the authors used random search of HPs, along with systematic removal of pieces of the LSTM architecture to determine which pieces of the LSTM were more important than the others. Our UGRNN architecture is directly inspired by the large impact of removing the forget gate from the LSTM (Gers et al., 1999). An in-depth comparison between RNNs and GRUs in the context of end-to-end speech recognition and a limited computational budget was conducted in Amodei et al. (2015). Further, ideas from RNN architectures that improve ease of training, such as forget gates (Gers et al., 1999), and copying recurrent state from one time step to another, are making their way into deep feed-forward networks as highway networks (Srivastava et al., 2015) and residual connections (He et al., 2015), respectively. Indeed, the +RNN was inspired in part by the coupled depth gate of Srivastava et al. (2015).

1.4 RECURRENT NEURAL NETWORK ARCHITECTURES

Below we briefly define the RNN architectures used in this study. Unless otherwise stated \mathbf{W} denotes a matrix, \mathbf{b} denotes a vector of biases. The symbol \mathbf{x}_t is the input at time t , \mathbf{h}_t is the hidden state at time t . Remaining vector variables represent intermediate values. The function $\sigma(\cdot)$ denotes the logistic sigmoid function and $s(\cdot)$ is either tanh or ReLU, set as a HP (see Appendix, Section RNN HPs for the complete list of HPs). Initial conditions for the networks were set to a learned bias. Finally, it is a well-known trick of the trade to initialize the gates of an LSTM or GRU with a large bias to induce better gradient flow. We included this parameter, denoted as b^{fg} , and tuned it along with all other HPs.

RNN, IRNN (LE ET AL., 2015)

$$\mathbf{h}_t = s(\mathbf{W}^h \mathbf{h}_{t-1} + \mathbf{W}^x \mathbf{x}_t + \mathbf{b}^h) \quad (1)$$

Note the IRNN is identical in structure to the vanilla RNN, but with an identity initialization for \mathbf{W}^h , zero initialization for the biases, and $s = \text{ReLU}$ only.

UGRNN - UPDATE GATE RNN

Based on Greff et al. (2015), where they noticed the forget gate “was crucial” to LSTM performance, we tried an RNN variant where we began with a vanilla RNN and *added* a single gate. This gate determines whether the hidden state is carried over from the previous time step, or updated – hence, it is an update gate. An alternative way to view the UGRNN is a highway layer gated through time (Srivastava et al., 2015).

$$\mathbf{c}_t = s(\mathbf{W}^{ch} \mathbf{h}_{t-1} + \mathbf{W}^{cx} \mathbf{x}_t + \mathbf{b}^c) \quad (2)$$

$$\mathbf{g}_t = \sigma(\mathbf{W}^{gh} \mathbf{h}_{t-1} + \mathbf{W}^{gx} \mathbf{x}_t + \mathbf{b}^g + b^{fg}) \quad (3)$$

$$\mathbf{h}_t = \mathbf{g}_t \cdot \mathbf{h}_{t-1} + (\mathbf{1} - \mathbf{g}_t) \cdot \mathbf{c}_t \quad (4)$$

GRU - GATED RECURRENT UNIT (CHO ET AL., 2014)

$$\mathbf{r}_t = \sigma(\mathbf{W}^{\text{rh}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{rx}} \mathbf{x}_t + \mathbf{b}^{\text{r}}) \quad (5)$$

$$\mathbf{u}_t = \sigma(\mathbf{W}^{\text{uh}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{ux}} \mathbf{x}_t + \mathbf{b}^{\text{u}} + b^{\text{fg}}) \quad (6)$$

$$\mathbf{c}_t = \text{s}(\mathbf{W}^{\text{ch}} (\mathbf{r}_t \cdot \mathbf{h}_{t-1}) + \mathbf{W}^{\text{cx}} \mathbf{x}_t + \mathbf{b}^{\text{c}}) \quad (7)$$

$$\mathbf{h}_t = \mathbf{u}_t \cdot \mathbf{h}_{t-1} + (1 - \mathbf{u}_t) \cdot \mathbf{c}_t \quad (8)$$

LSTM - LONG SHORT TERM MEMORY (HOCHREITER & SCHMIDHUBER, 1997)

$$\mathbf{i}_t = \sigma(\mathbf{W}^{\text{ih}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{ix}} \mathbf{x}_t + \mathbf{b}^{\text{i}}) \quad (9)$$

$$\mathbf{f}_t = \sigma(\mathbf{W}^{\text{fh}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{fx}} \mathbf{x}_t + \mathbf{b}^{\text{f}} + b^{\text{fg}}) \quad (10)$$

$$\mathbf{c}_t^{\text{in}} = \text{s}(\mathbf{W}^{\text{ch}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{cx}} \mathbf{x}_t + \mathbf{b}^{\text{c}}) \quad (11)$$

$$\mathbf{c}_t = \mathbf{f}_t \cdot \mathbf{c}_{t-1} + \mathbf{i}_t \cdot \mathbf{c}_t^{\text{in}} \quad (12)$$

$$\mathbf{o}_t = \sigma(\mathbf{W}^{\text{oh}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{ox}} \mathbf{x}_t + \mathbf{b}^{\text{o}}) \quad (13)$$

$$\mathbf{h}_t = \mathbf{o}_t \cdot \tanh(\mathbf{c}_t) \quad (14)$$

+RNN - INTERSECTION RNN

Due to the success of the UGRNN for shallower architectures in this study (see later figures on trainability), as well as some of the observed trainability problems for both the LSTM and GRU for deeper architectures (e.g. Figure 4h) we developed the Intersection RNN (denoted with a '+') architecture with a coupled depth gate in addition to a coupled recurrent gate. Additional influences for this architecture were the recurrent gating of the LSTM and GRU, and the depth gating from the highway network (Srivastava et al., 2015). This architecture has recurrent input, \mathbf{h}_{t-1} , and depth input, \mathbf{x}_t . It also has recurrent output, \mathbf{h}_t , and depth output, \mathbf{y}_t . Note that this architecture only applies between layers where \mathbf{x}_t and \mathbf{y}_t have the same dimension, and is not appropriate for networks with a depth of 1 (we exclude depth one +RNNs in our experiments).

$$\mathbf{y}_t^{\text{in}} = \text{s1}(\mathbf{W}^{\text{yh}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{yx}} \mathbf{x}_t + \mathbf{b}^{\text{y}}) \quad (15)$$

$$\mathbf{h}_t^{\text{in}} = \text{s2}(\mathbf{W}^{\text{hh}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{hx}} \mathbf{x}_t + \mathbf{b}^{\text{h}}) \quad (16)$$

$$\mathbf{g}_t^{\text{y}} = \sigma(\mathbf{W}^{\text{gyh}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{gyx}} \mathbf{x}_t + \mathbf{b}^{\text{gy}} + b^{\text{fg,y}}) \quad (17)$$

$$\mathbf{g}_t^{\text{h}} = \sigma(\mathbf{W}^{\text{ghh}} \mathbf{h}_{t-1} + \mathbf{W}^{\text{ghx}} \mathbf{x}_t + \mathbf{b}^{\text{gh}} + b^{\text{fg,h}}) \quad (18)$$

$$\mathbf{y}_t = \mathbf{g}_t^{\text{y}} \cdot \mathbf{x}_t + (1 - \mathbf{g}_t^{\text{y}}) \cdot \mathbf{y}_t^{\text{in}} \quad (19)$$

$$\mathbf{h}_t = \mathbf{g}_t^{\text{h}} \cdot \mathbf{h}_{t-1} + (1 - \mathbf{g}_t^{\text{h}}) \cdot \mathbf{h}_t^{\text{in}} \quad (20)$$

In practice we used ReLU for s1 and tanh for s2.

2 CAPACITY EXPERIMENTS

2.1 PER-PARAMETER CAPACITY

A foundational result in machine learning is that a single-layer perceptron with N^2 parameters can store at least 2 bits of information per parameter (Cover, 1965; Gardner, 1988; Baldi & Venkatesh, 1987). More precisely, a perceptron can implement a mapping from $2N$, N -dimensional, input vectors to arbitrary N -dimensional binary output vectors, subject only to the extremely weak restriction that the input vectors be in general position. RNNs provide a far more complex input-output mapping, with hidden units, recurrent dynamics, and a diversity of nonlinearities. Nonetheless, we wondered if there were analogous capacity results for RNNs that we might be able to observe empirically.

2.1.1 EXPERIMENTAL SETUP

As we will show in Section 3, tasks with complex temporal dynamics, such as language modeling, exhibit a per-parameter capacity bottleneck that explains the performance of RNNs far better than a

per-unit bottleneck. To make the experimental design as simple as possible, and to remove potential confounds stemming from the choice of temporal dynamics, we study per-parameter capacity using a task inspired by Gardner (1988). Specifically, to measure how much task-related information can be stored in the parameters of an RNN, we use a memorization task, where a random static input is injected into an RNN at a single time step, and a random static output is read out some number of time steps later. We emphasize that the same per-parameter bottleneck that we find in this simplified task also arises in more temporally complex tasks, such as language modeling.

At a high level, we draw a fixed set of random inputs and random labels, and train the RNN to map random inputs to randomly chosen labels via cross-entropy error. However, rather than returning the cross-entropy error to the HP tuner (as is normally done), we instead return the mutual information between the RNN outputs and the true labels. In this way, we can treat the number of input-output mappings as a HP, and the tuner will select for us the correct number of mappings so as to maximize the mutual information between the RNN outputs and the labels. From this mutual information we compute bits per parameter, which provides a normalized measurement of how much the RNN learned about the task.

More precisely, we draw datasets of binary inputs \mathbf{X} and target binary labels \mathbf{Y} at uniform from the set of all binary datasets, $\mathbf{X} \sim \mathcal{X} = \{0, 1\}^{n_{in} \times b}$, $\mathbf{Y} \sim \mathcal{Y} = \{0, 1\}^{1 \times b}$, where b is the number of samples, and n_{in} is the dimensionality of the inputs. As stated previously, b is treated as a HP. For each value of b the RNN is trained to minimize the cross entropy of the network output with the true labels. We write the output of the RNN for all inputs as $\hat{\mathbf{Y}} = f(\mathbf{X})$, with corresponding random variable $\hat{\mathcal{Y}}$. We are interested in the mutual information $I(\mathcal{Y}; \hat{\mathcal{Y}})$ between the true class labels and the class labels predicted by the RNN. This is the amount of (directly recoverable) information that the RNN has stored about the task. In this setting, it is calculated as

$$I(\mathcal{Y}; \hat{\mathcal{Y}}) = H(\mathcal{Y}) - H(\mathcal{Y}|\hat{\mathcal{Y}}) \quad (21)$$

$$= b + b(p \log_2 p + (1-p) \log_2 (1-p)), \quad (22)$$

where p is the fraction of correctly classified samples. The number b is then adjusted, along with all the other HPs, so as to maximize the mutual information $I(\mathcal{Y}; \hat{\mathcal{Y}})$. In practice p is computed using only a single draw of $\{\mathbf{X}, \mathbf{Y}\} \sim \mathcal{X} \times \mathcal{Y}$.

We performed this optimization of $I(\mathcal{Y}; \hat{\mathcal{Y}})$ for various RNN architectures, depths, and numbers of parameters. We plot the best value of $I(\mathcal{Y}; \hat{\mathcal{Y}})$ vs. number of parameters in Figure 1a. This captures the amount of information stored in the parameters about the mapping between \mathbf{X} and \mathbf{Y} . To get an estimate of bits per parameter, we divide by the number of parameters, as shown in Figure 1e.

2.1.2 RESULTS

Five Bits per Parameter Examining the results of Figure 1, we find the capacity of all architectures is roughly linear in the number of parameters, across several orders of magnitude of parameter count. We further find that the capacity is between 3 and 6 bits per parameter, once again across all architectures, depths 1, 2 and 4, and across several orders of magnitude in terms of number of parameters. Given the possibility of small size effects, and a larger portion of weights used as biases at a small number of parameters, we believe our estimates for larger networks are more reliable. This leads us to a bits per parameter estimate of approximately 5, averaging over all architectures and all depths. Finally, we note that the per-parameter task capacity increases as a function of the number of unrollings, though with diminishing gains (Figure 2b).

These consistent capacity results across diverse architectures and scales are even more surprising, since prior to these experiments it was not clear that capacity would even scale linearly with the number of parameters. For instance, previous results on model compression – by reducing the number of parameters (Yang et al., 2015), or by reducing the bit depth of parameters (Hubara et al., 2016) – might lead one to predict that different architectures use parameters with vastly different efficiencies, and that task capacity increases only sublinearly with parameter count.

Gating Slightly Reduces Capacity While overall, the different architectures performed very similarly, there are some capacity differences between architectures that appear to hold up across most depths and parameter counts. To quantify these differences we constructed a table showing the change in the number of parameters one would need to switch from one architecture to another, while maintaining equivalent capacity (Figure 1i). One trend that emerged from our capacity experiments is a slightly reduced capacity as a function of "gatedness". Putting aside the IRNN, which performed the worst and is discussed below, we noticed that across all depths and all model sizes, the performance was on average $\text{RNN} > \text{UGRNN} > \text{GRU} > \text{LSTM} > +\text{RNN}$. The vanilla RNN has no gates, the UGRNN has one, while the remaining three have two or more.

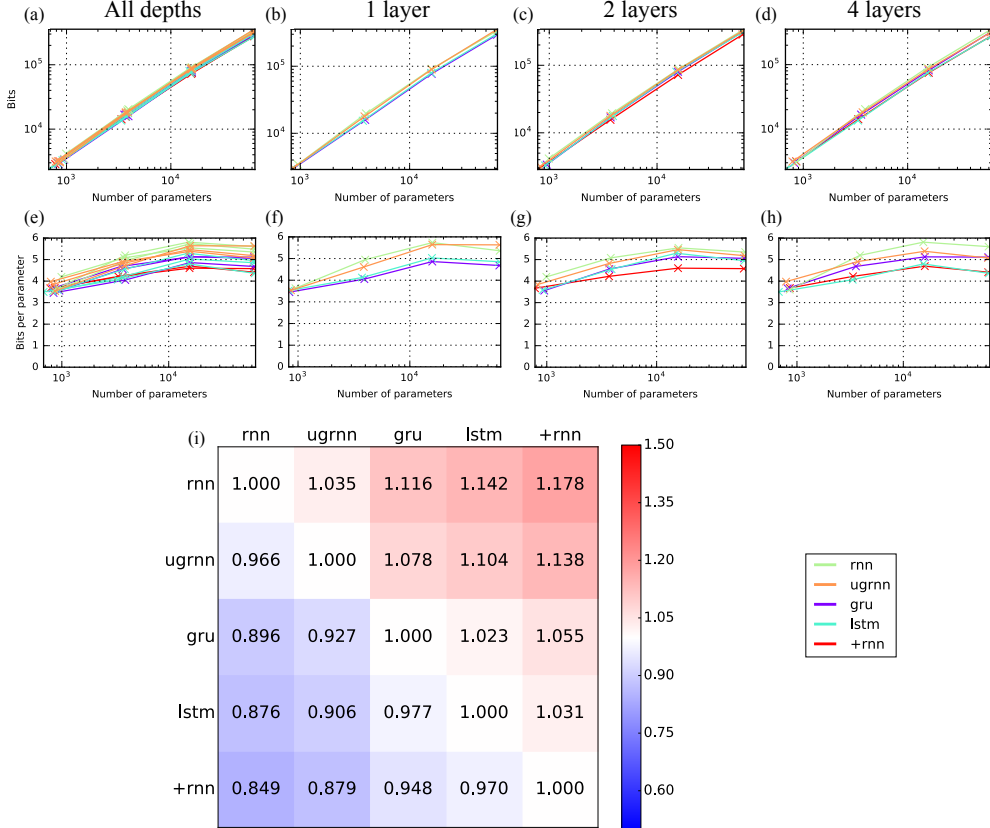


Figure 1: All neural network architectures can store approximately five bits per parameter about a task, with only small variations across architectures. (a) Stored bits as a function of network size. These numbers represent the maximum stored bits across 1000+ HP optimizations with 5 time steps unrolled at each network size for all levels of depth. (b-d) Same as (a), but each level of depth shown separately. (e-h) Same as (a-d) but showing bits per parameter as a function of network size. (i) The value in cell (x, y) is the multiplier for the number of parameters needed to give the architecture on the x -axis the same capacity as the architecture on the y -axis. Capacities measured by averaging the maximum stored bits per parameter for each architecture across all sizes and levels of depth.

ReLU's Reduce Capacity In our capacity tasks, the IRNN performed noticeably worse than all other architectures, reaching a maximum bits per parameter of roughly 3.5. To determine if this performance drop was due to the ReLU nonlinearity of the IRNN, or its identity initialization, we sorted through the RNN and UGRNN results (which both have ReLU and tanh as choices for the nonlinearity HP) and looked at the maximum bits per parameter when only optimizations using ReLU are considered. Indeed, both the RNN and UGRNN bits per parameter dropped dramatically to the 3.5 range (Figure 2a) when those architectures exclusively used ReLU, providing strong evidence that the ReLU activation function is problematic for this capacity task.

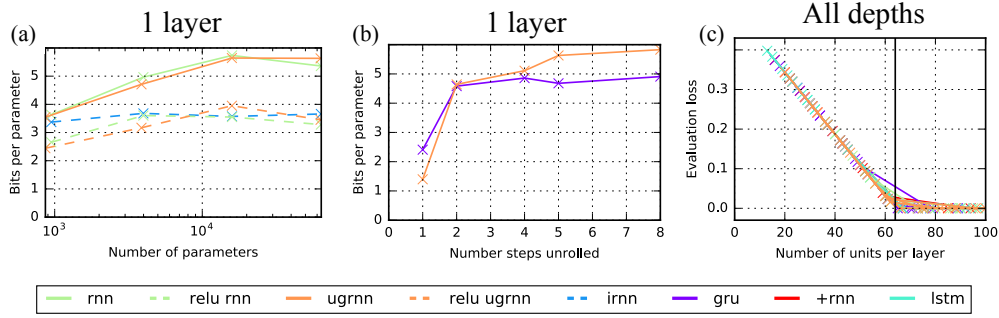


Figure 2: Additional RNN capacity analysis. (a) The effect of the ReLU nonlinearity on capacity. Solid lines indicate bits per parameter for 1-layer architectures (same as Figure 1b), where both tanh and ReLU are nonlinearity choices for the HP tuner. Dashed lines show the maximum bits per parameter for each architecture when only results achieved by the ReLU nonlinearity are considered. (b) Bits per parameter as a function of the number of time steps unrolled. (c) L2 error curve for all architectures of all depths on the memory throughput task. The curve shows the error plotted as a function of the number of units for a random input of dimension 64 (black vertical line). All networks with with less than 64 units have error in reconstruction, while all networks with number of units greater than 64 nearly perfectly reconstruct the random input.

2.2 PER-UNIT CAPACITY TO REMEMBER INPUTS

An additional capacity bottleneck in RNNs is their ability to store information about their inputs over time. It may be plainly obvious that an IRNN, which is essentially an integrator, can achieve perfect memory of its inputs if the number of inputs is less than or equal to the number of hidden units, but it is not so clear for some of the more complex architectures. So we measured the per-unit input memory empirically. Figure 2c, shows the intuitive result that every RNN architecture (at every depth and number of parameters) we studied can reconstruct a random n_{in} dimensional input at some time in the future, if and only if the number of hidden units per layer in the network, n_h , is greater than or equal to n_{in} . Moreover, regardless of RNN architecture, the error in reconstructing the input follows the same curve as a function of the number of hidden units for all RNN variants, corresponding to reconstructing an n_h dimensional subspace of the n_{in} dimensional input.

We highlight this per-unit capacity to make the point that a per-parameter task capacity appears to be the limiting factor in our experiments (e.g. Figure 1 and Figure 3), and not a per-unit capacity, such as the per-unit capacity to remember previous inputs. Thus when comparing results between architectures, one should normalize different architectures by the number of parameters, and not the number of units, as is frequently done in the literature (e.g. when comparing vanilla RNNs to LSTMs). This makes further sense as, for all common RNN architectures, the computational cost of processing a single sample is linear in the number of parameters, and quadratic in the number of units per layer. As we show in Figure 3d, plotting the capacity results by numbers of units gives very misleading results.

3 ADDITIONAL TASKS WHERE ARCHITECTURES ACHIEVE VERY SIMILAR LOSS

We studied additional tasks that we believed to be easy enough to train that the evaluation loss of different architectures would reveal variations in capacity rather than trainability. A critical aspect of these tasks is that they could not be learned perfectly by any of the model sizes in our experiments. As we change model size, we therefore expect performance on the task to also change. The tasks are (see Appendix, section Task Definitions for further elaboration of these tasks):

- text8 - 1-step ahead character-based prediction on the text8 Wikipedia dataset (100 million characters) (Mahoney, 2011).
- Random Continuous Functions (RCF) - A task similar to the per-parameter capacity task above, except the target outputs are real numbers (not categorical), and the number of training samples is held fixed.

The performance on these two tasks is shown in Figure 3. The evaluation loss as a function of the number of parameters is plotted in panels a-c and e-g, for the text8 task, and RCF task, respectively. For all tasks in this section, the number of parameters rather than the number of units provided the bottleneck on performance, and all architectures performed extremely closely for the same number of parameters. By close performance we mean that, for one model to achieve the same loss as another the model, the number of parameters would have to be adjusted by only a small factor (exemplified in Figure 1i for the per-parameter capacity task).

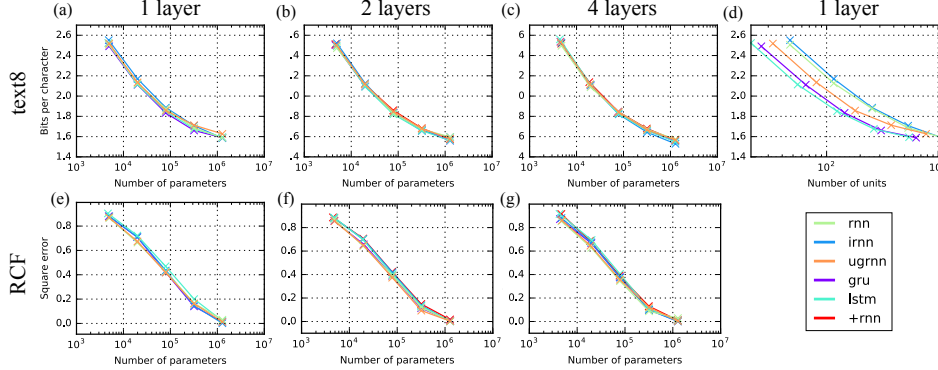


Figure 3: All RNN architectures achieved near identical performance given the same number of parameters, on a language modeling and random function fitting task. (a-c) text8 Wikipedia number of parameters vs bits per character for all RNN architectures. From left to right: 1 layer, 2 layer, 4 layer models. (d) text8 number of hidden units vs bits per character for 1 layer architectures. We note that this is almost always a misleading way to compare architectures. (e-g) Same as (a-c), except showing square error for different model sizes trained on RCFs.

4 TASKS THAT ARE VERY HARD TO LEARN

In practice it is widely appreciated that there is often a significant gap in performance between, for example, the LSTM and the vanilla RNN, with the LSTM nearly always outperforming the vanilla RNN. Our per-parameter capacity results provide evidence for a rough equivalence among a variety of RNN architectures, with slightly higher capacity in the vanilla RNN (Figure 1). To reconcile our per-parameter capacity results with widely held experience, we provide evidence that gated architectures, such as the LSTM, are far easier to train than the vanilla RNN (and often the IRNN).

We study two tasks that are difficult to learn: parallel parentheses counting of independent input streams, and mathematical addition of integers encoded in a character string (see Appendix, section Task Definitions). The parentheses task is moderately difficult to learn, while the arithmetic task is quite hard. The results of the HP optimizations are shown in Figure 4a-4h for the parentheses task, and in Figure 4i-4p for the arithmetic task. These tasks show that, while it is possible for a vanilla RNN to learn these tasks reasonably well, it is far more difficult than for a gated architecture. Note that the best achieved loss on the arithmetic task is still significantly decreasing, even after 2500 HP evaluations (2500 full complete optimizations over the training set), for the RNN and IRNN.

There are three noteworthy trends in these trainability experiments. First, across both tasks, and all depths (1, 2, 4 and 8), the RNN and IRNN performed most poorly, and took the longest to learn the task. Note, however that both the RNN and IRNN always solved the tasks eventually, at least for depth 1. Second, as the stacking depth increased, the gated architectures became the only architectures that could solve the tasks. Third, the most trainable architecture for depth 1 was the GRU, and the most trainable architecture for depth 8 was the +RNN.

To achieve our results on capacity and trainability, we relied heavily on a HP tuner. Most practitioners do not have the time or resources to make use of such a tuner, typically only adjusting the HPs a few times themselves. So we wondered how the various architectures would perform if we set HPs randomly, within the ranges specified (see Appendix for ranges). We tried this 1000 times on the parentheses task, for all 200k parameter architectures at depths 1 and 8 (Figure 5 and Table 1). The

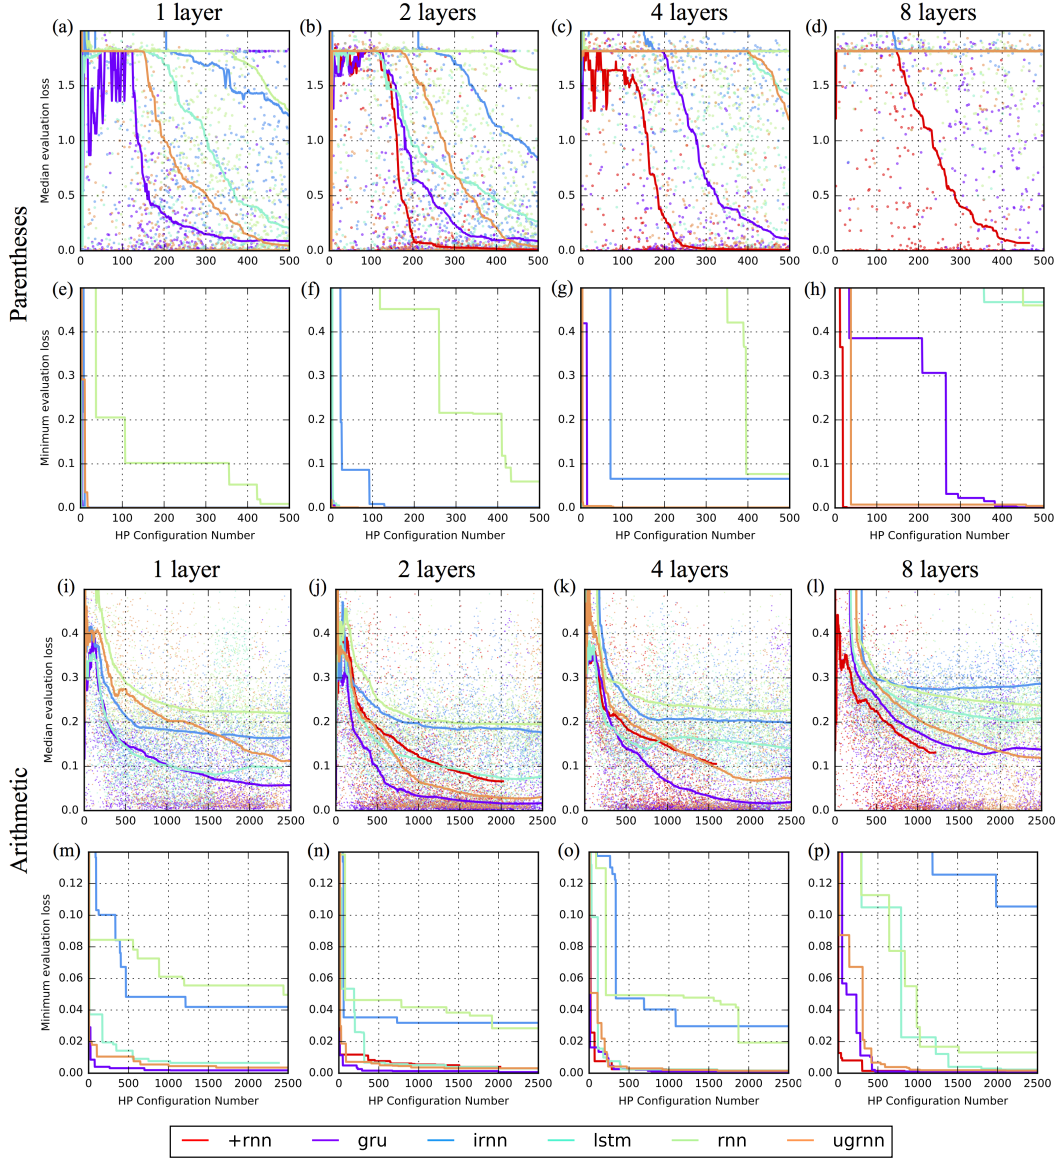


Figure 4: Some RNN architectures are far easier to train than others. Results of HP searches on extremely difficult tasks for the vanilla RNN. (a) Median evaluation error as a function of HP optimization iteration for 1 layer architectures on the parentheses task. Dots indicate evaluation loss achieved on that iteration. (b-d) Same as (a), but for 2, 4 and 8 layer architectures. (e-h) Minimum evaluation error as a function of HP optimization iteration for parentheses task. Same depth order as (a-d). (i-p) Same as (a-h), except for the arithmetic task. We note that the best loss for the vanilla RNN is still decreasing after 1500+ HP evaluations.

noticeable trends are that the IRNN returned an infeasible error nearly half of the time, and the LSTM (depth 1) and GRU (depth 8) were infeasible the least number of times, where infeasibility means that the training loss diverged. For depth 1, the GRU gave the smallest error, and the smallest median error, and for depth 8, the +RNN delivered the smallest error and smallest median error.




Figure 5: For generic hyperparameters, GRU and +RNN are the most easily trainable architectures. Evaluation losses from 1000 iterations of randomly chosen HP sets for 1 and 8 layer, 200k parameter models on the parentheses task. (a) Box and whisker plot of evaluation losses for the 1 layer model. Colored boxes show the span from the first to the third quartile of the evaluation losses. Whiskers indicate the minimum and maximum losses. Medians are shown in yellow, and crosses indicate outliers. (b) Same as (a) but for 8 layers.

Architecture	% Infeasible
GRU	15.5 %
IRNN	56.7 %
LSTM	12.0 %
RNN	21.5 %
UGRNN	20.2 %

Architecture	% Infeasible
GRU	3.2 %
IRNN	44.6 %
LSTM	4.0 %
RNN	18.7 %
UGRNN	11.5 %
+RNN	8.8 %

Table 1: Fraction infeasible trials as a result of 1000 iterations of randomly chosen HP sets for 1 layer (left) and 8 layer (right), 200k parameter models on the parentheses task.

5 DISCUSSION

Here we report that a number of RNN variants can hold between 3-6 bits per parameter about their task, and that these variants can remember a number of random inputs that is nearly equal to the number of hidden units in the RNN. The quantification of the number of bits per parameter an RNN can store about a task is particularly important, as it was not previously known whether the amount of information about a task that could be stored was even linear in the number of parameters.

While our results point to empirical capacity limits for both task memorization, and input memorization, apparently the requirement to remember features of the input through time is not a practical bottleneck. If it were, then the vanilla RNN and IRNN would perform better than the gated architectures in proportion to the ratio of the number of units, which they do not. Based on widespread results in the literature, and our own results on our difficult tasks, the loss of some memory capacity (and possibly a small amount of per-parameter storage capacity) for improved trainability seems a worthwhile trade off. Indeed, the input memory capacity did not obviously impact any task not explicitly designed to measure it, as the error curves – for instance for the language modeling task – overlapped across architectures for the same number of parameters, but not the same number of units.

Our result on per-parameter task capacity, about 5 bits per parameter averaged over architectures, is in surprising agreement with recently published results on the capacity of synapses in biological neurons. This number was recently calculated to be about 4.7 bits per synapse, based on biological synapses in the hippocampus having roughly 26 measurable discrete sizes (Bartol et al., 2016). Our capacity results have implications for compressed networks that employ quantization techniques. In particular, they provide an estimate of the number of bits which a weight may be compressed without loss in task performance. Coincidentally, in Han et al. (2015), the authors used 5 bits per weight in the fully connected layers.

An additional observation about per-parameter task capacity in our experiments is that it increases for a few time steps beyond one (Figure 2b), and then appears to saturate. We interpret this to suggest that recurrence endows additional capacity to a network with shared parameters, but that there are diminishing returns, and the total capacity remains bounded even as the number of time steps increases.

Despite our best efforts, we cannot claim that we perfectly trained any of the models. Potential problems in HP optimization could be local minima, as well as stochastic behavior in the HP optimization as a result of the stochasticity of batching or random draws for weight matrices. We tried to uncover these effects by running the best performing HPs 100 times, and did not observe any serious deviations from the best results (see Table App.1 in Appendix). Another form of validation comes from the fact that in our capacity task, essentially 3 independent experiments (one for each level of depth) yielded a clustering by architecture (Figure 1e).

Do our results yield a framework for choosing a recurrent architecture? In total, we believe yes. As explored in Amodei et al. (2015), a practical concern for recurrent models is speed of execution in a production environment. Our results suggest that if one has a large resource budget for training and confined resource budget for inference, one should choose the vanilla RNN. Conversely, if the training resource budget is small, but the inference budget large, one should choose a gated model. Another serious concern relates to task complexity. If the task is easy to learn, a vanilla RNN should yield good results. However if the task is even moderately difficult to learn, a gated architecture is the right choice. Our results point to the GRU as being the most learnable of gated architectures for shallow architectures, followed by the UGRNN. The +RNN typically performed best for deeper architectures. Our results on trainability confirm the widely held view that the LSTM is an extremely reliable architecture, but it was almost never the best performer in our experiments. All things considered, in an uncertain training environment, we would recommend using the GRU or +RNN.

REFERENCES

- Dario Amodei, Rishita Anubhai, Eric Battenberg, Carl Case, Jared Casper, Bryan C. Catanzaro, Jingdong Chen, Mike Chrzanowski, Adam Coates, Greg Diamos, Erich Elsen, Jesse Engel, Linxi Fan, Christopher Fougner, Tony Han, Awni Y. Hannun, Billy Jun, Patrick LeGresley, Libby Lin, Sharan Narang, Andrew Y. Ng, Sherjil Ozair, Ryan Prenger, Jonathan Raiman, Sanjeev Satheesh, David Seetapun, Shubho Sengupta, Yi Wang, Zhiqian Wang, Chong Wang, Bo Xiao, Dani Yogatama, Jun Zhan, and Zhenyao Zhu. Deep speech 2: End-to-end speech recognition in english and mandarin. *CoRR*, abs/1512.02595, 2015. URL <http://arxiv.org/abs/1512.02595>.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Pierre Baldi and Santosh S Venkatesh. Number of stable points for spin-glasses and neural networks of higher orders. *Physical Review Letters*, 58(9):913, 1987.
- Thomas M Bartol, Cailey Bromer, Justin Kinney, Michael A Chirillo, Jennifer N Bourne, Kristen M Harris, and Terrence J Sejnowski. Nanocconnectomic upper bound on the variability of synaptic plasticity. *eLife*, 4: e10778, 2016.
- Adam S Charles, Han Lun Yap, and Christopher J Rozell. Short-term memory capacity in networks via the restricted isometry property. *Neural computation*, 26(6):1198–1235, 2014.
- Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- Thomas M Cover. Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE transactions on electronic computers*, (3):326–334, 1965.
- Thomas Desautels, Andreas Krause, and Joel W Burdick. Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization. *The Journal of Machine Learning Research*, 15(1):3873–3923, 2014.
- Kenji Doya. Universality of fully connected recurrent neural networks. *Dept. of Biology, UCSD, Tech. Rep.*, 1993.

- Jakob Foerster, Justin Gilmer, Jan Chorowski, Jascha Sohl-Dickstein, and David Sussillo. Intelligible language modeling with input switched affine networks. *ICLR 2017 submission*, 2016.
- Surya Ganguli, Dongsung Huh, and Haim Sompolsky. Memory traces in dynamical systems. *Proceedings of the National Academy of Sciences*, 105(48):18970–18975, 2008.
- Elizabeth Gardner. The space of interactions in neural network models. *Journal of physics A: Mathematical and general*, 21(1):257, 1988.
- Felix A. Gers, Jurgen Schmidhuber, and Fred Cummins. Learning to forget: Continual prediction with lstm. *Artificial Neural Networks, ICANN 99. Ninth International Conference on (Conf. Publ. No. 470)*, 2:850–855, 1999.
- Alex Graves, Marcus Liwicki, Santiago Fernández, Roman Bertolami, Horst Bunke, and Jürgen Schmidhuber. A novel connectionist system for unconstrained handwriting recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 31(5):855–868, 2009.
- Klaus Greff, Rupesh Kumar Srivastava, Jan Koutník, Bas R Steunebrink, and Jürgen Schmidhuber. Lstm: A search space odyssey. *arXiv preprint arXiv:1503.04069*, 2015.
- Karol Gregor, Ivo Danihelka, Alex Graves, and Daan Wierstra. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015.
- Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Itay Hubara, Daniel Soudry, and Ran El Yaniv. Binarized neural networks. *arXiv preprint arXiv:1602.02505*, 2016.
- Herbert Jaeger and Harald Haas. Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. *science*, 304(5667):78–80, 2004.
- Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever. An empirical exploration of recurrent network architectures. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pp. 2342–2350, 2015.
- Andrej Karpathy, Justin Johnson, and Fei-Fei Li. Visualizing and understanding recurrent networks. *arXiv preprint arXiv:1506.02078*, 2015.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. URL <http://arxiv.org/abs/1412.6980>.
- Pascal Koiran and Eduardo D Sontag. Vapnik-chervonenkis dimension of recurrent neural networks. *Discrete Applied Mathematics*, 86(1):63–79, 1998.
- Quoc V Le, Navdeep Jaitly, and Geoffrey E Hinton. A simple way to initialize recurrent networks of rectified linear units. *arXiv preprint arXiv:1504.00941*, 2015.
- Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, 14(11):2531–2560, 2002.
- Matt Mahoney. Large text compression benchmark: About the test data, 2011. URL <http://mattmahoney.net/dc/textdata>. [Online; accessed 15-November-2016].
- Valerio Mante, David Sussillo, Krishna V Shenoy, and William T Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84, 2013.
- James Martens and Ilya Sutskever. Learning recurrent neural networks with hessian-free optimization. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 1033–1040, 2011.
- Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein. Deep knowledge tracing. In *Advances in Neural Information Processing Systems*, pp. 505–513, 2015.

- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems*, pp. 2951–2959, 2012.
- Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. Highway networks. *arXiv preprint arXiv:1505.00387*, 2015.
- David Sussillo and Omri Barak. Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural computation*, 25(3):626–649, 2013.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pp. 3104–3112, 2014.
- Tijmen Tieleman and Geoffrey. Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4, 2012.
- Olivia L White, Daniel D Lee, and Haim Sompolinsky. Short-term memory in orthogonal neural networks. *Physical review letters*, 92(14):148102, 2004.
- Zichao Yang, Marcin Moczulski, Misha Denil, Nando de Freitas, Alex Smola, Le Song, and Ziyu Wang. Deep fried convnets. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1476–1483, 2015.

Appendix

A RNN HPs SET BY THE HP TUNER

We used a HP tuner that uses a Gaussian Process (GP) Bandits approach for HP optimization. Our setting of the tuner’s internal parameters was such that it uses Batched GP Bandits with an expected improvement acquisition function and a Matern 5/2 Kernel with feature scaling and automatic relevance determination performed by optimizing over kernel HPs. Please see Desautels et al. (2014) and Snoek et al. (2012) for closely related work.

For all our tasks, we requested HPs from the tuner, and reported loss on a validation dataset. For the per-parameter capacity task, the validation and training datasets were identical. For all other tasks the validation set consisted of held out data. Additionally, the performance we plot in all cases is on an evaluation dataset that is separate from both training and validation datasets.

Below is the list of all tunable HPs that were generically applied to all models. In total, each RNN variant had between 10 and 27 HP dimensions relating to the architecture, optimization, and regularization.

- $s()$ - as used in the following RNN definitions, a nonlinearity determined by the HP tuner, $\{\text{ReLU}, \tanh\}$. The only exception was the IRNN, which used ReLU exclusively.
- For any matrix that is inherently square, e.g. \mathbf{W}^h , there were three possible initializations: identity, orthogonal, or random normal distribution scaled by $1/\sqrt{n_h}$, with n_h the number of recurrent units. The sole exception was the RNN, which was limited to either orthogonal or random normal initializations, to differentiate it from the IRNN. For any matrix that is inherently rectangular, e.g. \mathbf{W}^x , we initialized with a random normal distribution scaled by $1/\sqrt{n_{in}}$, with n_{in} the number of inputs.
- For all matrix initializations except the identity initialization, there was a multiplicative scalar used to set the scale of matrix. The scalar was exponentially distributed in $[0.01, 2.0]$ for recurrent matrices and $[0.001, 2.0]$ for rectangular matrices.
- Biases could have two possible distributions: all biases set to a constant value, or drawn from a standard normal distribution.
- For all bias initializations, a multiplicative scalar was drawn, uniformly distributed in $[-2.0, 2.0]$ and applied to bias initialization.
- We included a scalar bias HP b^{fg} for architectures that contain forget or update gates, as is commonly employed in practice, which was uniformly distributed in $[0.0, 6.0]$.

Additionally, the HP tuner was used to optimize HPs associated with learning:

- The number of training steps - The exact range varied between tasks, but always fell between 50K and 20M.
- One of four optimization algorithms could be chosen: vanilla SGD, SGD with momentum, RMSProp (Tieleman & Hinton, 2012), or ADAM (Kingma & Ba, 2014).
- learning rate initial value, exponentially distributed in $[1e-4, 1e-1]$
- learning rate decay - exponentially distributed in $[1e-3, 1]$. The learning rate exponentially decays by this factor over the number of training steps chosen by the tuner
- optimizer momentum-like parameter - expressed as a logit, and uniformly distributed in $[1.0, 7.0]$
- gradient clipping value - exponentially distributed in $[1, 100]$
- l2 decay - exponentially distributed in $[1e-8, 1e-3]$.

Some optimization algorithms had additional parameters such as ADAM’s second order decay rate, or epsilon parameter. These were set to their default values and not optimized. We note also that the batch size was set individually by hand for all experiments, and the same random seed was used to initialize all network and task parameters. Note that for each network, the initial condition was set to a learned vector.

B TASK DEFINITIONS

PERCEPTRON CAPACITY

While at a high-level, for the perceptron capacity task, we wanted to optimize the amount of information the RNN carried about true random labels, in practice, the training objective was standard cross-entropy. However, when returning a validation loss to the HP tuner, we returned the mutual information $I(\mathbf{Y}; \hat{\mathbf{Y}}|\mathbf{X})$. Conceptually, this is as if there is one nested optimization inside another. The inner loop optimizes the RNN for the set of HPs, training cross entropy, but returning mutual information. The outer loop then chooses the HPs, in particular, the number of samples b , in equation (21), so as to maximize the amount of mutual information. This implementation is necessitated because there is no straightforward way to differentiate mutual information with respect to number of samples.

A HP determined whether the input vector \mathbf{x}_j was presented to the RNN only at the first time step, or whether it was presented at every time step.

MEMORY CAPACITY

In the Memory Capacity task, we wanted to know how much information an RNN can reconstruct about its inputs at some later time point. We picked an input dimension, 64, and varied the number of parameters in the networks such that the number of hidden units was roughly centered around 64. After 12 time steps the target of the network was exact reconstruction of the input, with a square error loss. The inputs were random values drawn from a uniform distribution.

RANDOM CONTINUOUS FUNCTION

A dataset was constructed consisting of $N = 10^6$ random unit norm Gaussian input vectors \mathbf{x} , with size $d = 50$. Target scalar outputs y were generated for each input vector, and were also drawn from a unit norm Gaussian. Each sample i was assigned a power law weighting $\beta_i = \frac{(i+\tau)^{-1}}{Z}$, where Z was a normalization constant such that the weightings summed to 1, and the characteristic time constant $\tau = 5000$. The loss function for training was weighted square error on the y_i , with the β_i acting as the weighting terms.

TEXT8

In the text8 task, the task was to predict one character ahead in the text8 dataset (1e8 characters of Wikipedia) (Mahoney, 2011). Input was a hot-one encoded sequence, as was the output. The loss was cross-entropy loss on a softmax output layer. Rather than use partial unrolling as is common in language modeling, we generated random pointers into the text. The first $T/4$ time steps were used to burn in the RNN, and remaining steps were used for training or inference.

PARENTHESES COUNTING TASK

The parentheses counting task independently counts the number of opened 'parens', e.g. '(', without the closing ')'. Here parens is used to mean any of 10 parens type pairs, e.g. '<>' or '[]'. Additionally, there were 10 noise characters, 'a' to 'j'. For each paren type, there was a $20D + 10D = 30D$ hot-one encoding of all paren and noise symbols, for a total of 300 inputs. The output for each paren type was a hot-one encoding of the digits 0-9, which represented the count of the opened parens of that type. If the count exceeded 9, the the network kept the count at 9, if the paren was closed, the count decreased. The loss was the sum of cross-entropy losses, one for each paren type. Finally, for each paren input stream, 50% random noise characters were drawn, and 50% random paren characters were drawn, e.g. 10 streams like '(a<a<bcb>[[[)'. Parens of other types were treated as noise for the current type, e.g. for the above string if the paren type was '<>', the answer is '1' at the end. The loss was defined only at the final time point, T , and $T = 175$.

ARITHMETIC TASK

In the arithmetic task, a hot-one encoded character sequence of an addition problem was presented as input to the network, e.g., '-343243+93851=', and the output was the hot-one encoded answer, including the correct amount of left padded spaces, '-249392'. An additional HP for this task was the number of compute steps (1-6) between the input of the '=' and the first non-space character in the target output sequence. The two numbers in the input were randomly, uniformly selected in $[-1e7, 1e7]$. We found this task to be extremely difficult for the networks to learn, but when the task was learned, certain of the network architectures could perform the task nearly perfectly.

HP ROBUSTNESS

We wondered how robust the HPs are to the variability of both random batching of data, and random initialization of parameters. So we identified the best HPs from the parentheses experiments of 100k parameter, 1 layer architectures, and reran the parameter optimization 100 times. We measured the number of infeasible experiments, as well as a number of statistics of the loss for the reruns (Table App.1). These results show that the best HPs yielded a distribution of losses very close to the originally reported loss value.

Architecture	Original	Infeasible	Min	Mean	Max	S.D.	S.D./Mean
RNN	1.16e-2	0 %	1.41e-2	8.21e-2	0.294	5.22e-2	0.636
IRNN	4.20e-4	48 %	2.24e-4	5.02e-4	8.69e-4	1.35e-4	0.269
UGRNN	1.02e-4	0 %	3.66e-5	2.71e-4	6.06e-3	7.12e-4	2.63
GRU	2.80e-4	1 %	7.66e-5	1.89e-4	5.48e-4	9.08e-5	0.480
LSTM	7.96e-4	0 %	8.10e-4	2.02e-3	0.0145	2.31e-3	1.14

Table App.1: Results of 100 runs on the parentheses task using the best HPs for each architecture, at depth 1. HPs were chosen to be the set which achieved the minimum loss. Table shows original loss achieved by the HP tuner, amount of infeasible trials, minimum loss from running 100 iterations of the same HPs, mean loss, maximum loss, standard deviation, and standard deviation divided by the mean.